arXiv:1402.2823v1 [math.ST] 12 Feb 2014

High-dimensional tests for spherical location and spiked covariance

Christophe Ley*, Davy Paindaveine† and Thomas Verdebout‡

Université Libre de Bruxelles and Université Lille Nord de France

Abstract

Rotationally symmetric distributions on the p-dimensional unit hypersphere, extremely popular in directional statistics, involve a location parameter $\boldsymbol{\theta}$ that indicates the direction of the symmetry axis. The most classical way of addressing the spherical location problem $\mathcal{H}_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$, with $\boldsymbol{\theta}_0$ a fixed location, is the so-called Watson test, which is based on the sample mean of the observations. This test enjoys many desirable properties, but its implementation requires the sample size n to be large compared to the dimension p. This is a severe limitation, since more and more problems nowadays involve high-dimensional directional data (e.g., in genetics or text mining). In this work, we therefore introduce a modified Watson statistic that can cope with high-dimensionality. We derive its asymptotic null distribution as both n and p go to infinity. This is achieved in a universal asymptotic framework that allows p to go to infinity arbitrarily fast (or slowly) as a function of n. We further show that our results also provide high-dimensional tests for a problem that has recently attracted much attention, namely that of testing that the covariance matrix of a multinormal distribution has a "$\boldsymbol{\theta}_0$-spiked" structure. Finally, a Monte Carlo simulation study corroborates our asymptotic results.

Keywords: Directional statistics, high-dimensional data, location tests, principal component analysis, rotationally symmetric distributions, spherical mean

1 Introduction

The technological advances and the ensuing new devices to collect and store data lead nowadays in many disciplines to data sets with very high dimension p, often larger than the sample size n. Consequently, there is a need for inferential methods that can deal with such high-dimensional data, and this has entailed a huge activity related to high-dimensional problems in the last decade. One- and multi-sample location problems have been investigated in [21], [20], [8], [22], and [23], among others. Since the seminal paper [13], problems related to covariance or scatter matrices have also been thoroughly studied by several authors; see, e.g., [9], [14], [16] and [11].

*E-mail address: chrisley@ulb.ac.be; URL: http://homepages.ulb.ac.be/~chrisley

†E-mail address: dpaindav@ulb.ac.be; URL: http://homepages.ulb.ac.be/~dpaindav

‡E-mail address: thomas.verdebout@univ-lille3.fr; URL: http://perso.univ-lille3.fr/~tverdebout

In this paper, we are interested in high-dimensional *directional* data, that is, in data lying on the unit hypersphere $\mathcal{S}^{p-1} = \{\mathbf{x} \in \mathbb{R}^p : \|\mathbf{x}\| = \sqrt{\mathbf{x}'\mathbf{x}} = 1\}$, with p large. Such data occur when only the direction of the observations and not their magnitude matters, and are extremely common, e.g., in magnetic resonance (10), gene-expression (1), and text mining (2). Inference for high-dimensional directional data has already been considered in several papers. For instance, [3, 4] and [2] investigate clustering methods in this context. Most asymptotic results from the literature, however, have been obtained as p goes to infinity, with n fixed. This is the case of almost all results in [24], [25], [27], and [10]. To the best of our knowledge, the only (n, p)-asymptotic results available can be found in [10], [7], [6], and [18]. However, [10] imposes the stringent condition that $p/n^2 \to \infty$ when studying the asymptotic behavior of the classical pseudo-FvML location estimator (FvML here refers to *Fisher-von Mises-Langevin* distributions; see below). [7] and [6] consider various (n, p)-asymptotic regimes in the context of testing for uniformity on the unit sphere, but the tests to be used depend on the regime considered which makes practical implementation problematic. Finally, [18] propose tests that are robust to the (n, p)-asymptotic regime considered; their tests, however, are sign procedures, hence are not based on sufficient statistics — unlike the much more classical pseudo-FvML procedures.

In the present paper, we intend to overcome these limitations in the context of the spherical location problem, one of the most fundamental problems in directional statistics. The natural distributional framework for this problem is provided by *rotationally symmetric distributions* (see Section 2), that form a semiparametric model, indexed by a finite-dimensional (location) parameter $\boldsymbol{\theta} \in \mathcal{S}^{p-1}$ and an infinite-dimensional parameter F. The spherical location problem consists in testing the null hypothesis $\mathcal{H}_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$ against alternative locations, where $\boldsymbol{\theta}_0$ is a given unit vector and F remains unspecified. The classical test for this problem is the so-called Watson test, based on the sample mean of the observations; see [26]. This test enjoys many desirable properties, and in particular is a *pseudo-FvML* procedure : in other words, it achieves optimality under FvML distributions, yet remains valid (in the sense that it meets the asymptotic nominal level constraint) under extremely mild assumptions on F.

Unfortunately, the Watson test cannot be used in the high-dimensional case, since its implementation crucially relies on fixed-p asymptotic results. In view of the growing number of high-dimensional directional data to analyze, this is a severe limitation. The aim of this paper hence is to define a modified Watson test statistic that can cope with high-dimensionality. We achieve this in such a way that asymptotic validity under virtually any rotationally symmetric distribution is maintained. Even better : in contrast with earlier asymptotic investigations of high-dimensional pseudo-FvML procedures, our asymptotic results are "universal" in the sense that they only require that p goes to infinity as n does (p may go arbitrarily fast (or slowly) to infinity as a function of n). Moreover, as a highly interesting by-product, we show that our procedure can be used to test the null hypothesis that the covariance matrix of a high-dimensional multinormal distribution is "$\boldsymbol{\theta}_0$-spiked", meaning that it is of the form $\boldsymbol{\Sigma} = \sigma^2(\mathbf{I}_p + \lambda\boldsymbol{\theta}_0\boldsymbol{\theta}_0')$ for some $\sigma^2, \lambda > 0$ and $\boldsymbol{\theta}_0 \in \mathbb{R}^k$; see, e.g., [12] or the quite recent [16] where this covariance structure has been used as an alternative to sphericity.

The outline of the paper is as follows. In Section 2, we define the class of rotationally symmetric distributions and introduce the Watson test for spherical location. In Section 3, we propose a modified Watson test statistic and derive its asymptotic null distribution in the high-

dimensional setting. We also prove that, in some cases, it is asymptotically equivalent to a sign test statistic. In Section 4, we show that the modified Watson test as well permits to test for a spiked covariance structure in multinormal distributions. A Monte Carlo simulation study is conducted in Section 5, while an Appendix collects the proofs of some technical lemmas.

2 Rotational symmetry and the Watson test

The distribution of the random p-vector $\mathbf{X}$, with values on the unit hypersphere $\mathcal{S}^{p-1}$, is *rotationally symmetric* about location $\boldsymbol{\theta}(\in \mathcal{S}^{p-1})$ if $\mathbf{OX}$ is equal in distribution to $\mathbf{X}$ for any orthogonal $p\times p$ matrix $\mathbf{O}$ satisfying $\mathbf{O}\boldsymbol{\theta}=\boldsymbol{\theta}$; see [19]. Rotationally symmetric distributions are characterized by the location parameter $\boldsymbol{\theta}$ and an infinite-dimensional parameter, the cumulative distribution function F of $\mathbf{X}'\boldsymbol{\theta}$, hence they are of a semiparametric nature. The rotationally symmetric distribution associated with $\boldsymbol{\theta}$ and F will be denoted as $\mathcal{R}(\boldsymbol{\theta},F)$ in the sequel. The most celebrated members of this family are the Fisher-von Mises-Langevin distributions, corresponding to $F_{p,\kappa}(t) = c_{p,\kappa}\int_{-1}^{t}(1-s^2)^{(p-3)/2}\exp(\kappa s)\,ds$ $(t\in[-1,1])$, where $c_{p,\kappa}$ is a normalization constant and $\kappa(>0)$ is a *concentration* parameter (the larger the value of κ, the more concentrated about $\boldsymbol{\theta}$ the distribution is); see [15] for further details.

Let $\mathbf{X}_1,\ldots,\mathbf{X}_n$ be a sequence of i.i.d. random unit vectors from $\mathcal{R}(\boldsymbol{\theta},F)$ and consider the problem of testing the null hypothesis $\mathcal{H}_0:\boldsymbol{\theta}=\boldsymbol{\theta}_0$ against the alternative $\mathcal{H}_1:\boldsymbol{\theta}\neq\boldsymbol{\theta}_0$, where $\boldsymbol{\theta}_0\in\mathcal{S}^{p-1}$ is fixed and F remains unspecified. Letting $\bar{\mathbf{X}}:=\frac{1}{n}\sum_{i=1}^{n}\mathbf{X}_i$, the classical test for this problem rejects the null for large values of the Watson statistic

$$W_n := \frac{n(p-1)\bar{\mathbf{X}}'(\mathbf{I}_k-\boldsymbol{\theta}_0\boldsymbol{\theta}_0')\bar{\mathbf{X}}}{1-\frac{1}{n}\sum_{i=1}^{n}(\mathbf{X}_i'\boldsymbol{\theta}_0)^2}. \tag{2.1}$$

Under very mild assumptions on F, the fixed-p asymptotic null distribution of W_n is chi-square with $p-1$ degrees of freedom. The resulting test, ϕ_n^W say, therefore rejects the null, at asymptotic level α, whenever $W_n>\Psi_{p-1}^{-1}(1-\alpha)$, where Ψ_{p-1} stands for the cumulative distribution function of the chi-square distribution with $p-1$ degrees of freedom; see [26].

Beyond achieving asymptotic level α under virtually any rotationally symmetric distribution, ϕ_n^W is optimal — more precisely, locally and asymptotically maximin, in the Le Cam sense — when the underlying distribution is FvML; for details, we refer to [17], where the asymptotic properties of ϕ_n^W under local alternatives are derived. Although ϕ_n^W is based on the sample mean of the observations, these excellent power properties are not obtained at the expense of robustness, since observations by construction are on the unit hypersphere.

Consequently, ϕ_n^W is a nice solution to the testing problem considered on all counts but one : implementation is based on fixed-p asymptotics, so that ϕ_n^W cannot be used when p is of the same order as, or even larger than, n. The goal of the present work is therefore to derive a modified Watson test, $\tilde{\phi}_n^W$ say, that can cope with high-dimensionality.

3 A high-dimensional Watson test

Consider the high-dimensional version of the testing problem $\mathcal{H}_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$ against $\mathcal{H}_1 : \boldsymbol{\theta} \neq \boldsymbol{\theta}_0$, based on a triangular array of observations $\mathbf{X}_{ni}$, $i = 1, \ldots, n$, $n = 1, 2, \ldots$, where $\mathbf{X}_{ni}$ takes values in $\mathcal{S}^{p_n - 1}$ and p_n goes to infinity with n. In this section, we modify the Watson test statistic W_n in (2.1) to make it robust to high-dimensionality. To do so, consider the (null) *tangent-normal decomposition* $\mathbf{X}_{ni} = (\mathbf{X}_{ni}'\boldsymbol{\theta}_0)\boldsymbol{\theta}_0 + u_{ni}\mathbf{S}_{ni}$, where

$$u_{ni} := \sqrt{1 - (\mathbf{X}_{ni}'\boldsymbol{\theta}_0)^2} \quad \text{and} \quad \mathbf{S}_{ni} := \frac{\mathbf{X}_{ni} - (\mathbf{X}_{ni}'\boldsymbol{\theta}_0)\boldsymbol{\theta}_0}{\|\mathbf{X}_{ni} - (\mathbf{X}_{ni}'\boldsymbol{\theta}_0)\boldsymbol{\theta}_0\|},$$

and note that the Watson statistic rewrites

$$\begin{aligned} W_n = \frac{p_n - 1}{\sum_{i=1}^n u_{ni}^2} \sum_{i,j=1}^n u_{ni}u_{nj}\mathbf{S}_{ni}'\mathbf{S}_{nj} &= \frac{p_n - 1}{\sum_{i=1}^n u_{ni}^2}\left(\sum_{i=1}^n u_{ni}^2 + 2\sum_{1\leq i<j\leq n} u_{ni}u_{nj}\mathbf{S}_{ni}'\mathbf{S}_{nj}\right) \\ &= (p_n - 1) + \frac{2(p_n - 1)}{\sum_{i=1}^n u_{ni}^2} \sum_{1\leq i<j\leq n} u_{ni}u_{nj}\mathbf{S}_{ni}'\mathbf{S}_{nj}. \end{aligned}$$

We then introduce the modified statistic

$$\tilde{W}_n := \frac{W_n - (p_n - 1)}{\sqrt{2(p_n - 1)}} = \left(\frac{\sqrt{2(p_n - 1)}}{n\mathrm{E}[u_{n1}^2]} \sum_{1\leq i<j\leq n} u_{ni}u_{nj}\mathbf{S}_{ni}'\mathbf{S}_{nj}\right) \Big/ \left(\frac{\frac{1}{n}\sum_{i=1}^n u_{ni}^2}{\mathrm{E}[u_{n1}^2]}\right). \tag{3.2}$$

The following result, that provides the (n, p)-asymptotic null distribution of $\tilde{W}_n$, is the main result of the paper.

Theorem 3.1. *Let $\mathbf{X}_{ni}$, $i = 1, \ldots, n$, $n = 1, 2, \ldots$, form a triangular array of random vectors satisfying the following conditions : (i) for any n, $\mathbf{X}_{n1}, \mathbf{X}_{n2}, \ldots, \mathbf{X}_{nn}$ are mutually independent and share a common rotationally symmetric distribution on $\mathcal{S}^{p_n - 1}$ with location $\boldsymbol{\theta}_0$; (ii) $p_n \to \infty$ as $n \to \infty$; (iii) $\mathrm{E}[u_{n1}^2] > 0$ and (iv) $\mathrm{E}[u_{n1}^4]/(\mathrm{E}[u_{n1}^2])^2 = o(n)$ as $n \to \infty$. Then $\tilde{W}_n$ is asymptotically standard normal.*

The assumptions of Theorem 3.1 are extremely mild. Note in particular that it is not assumed that the common distribution of the $\mathbf{X}_{ni}$'s is absolutely continuous with respect to the surface area measure on $\mathcal{S}^{p_n - 1}$. Imposing (iii) is strictly equivalent to requiring that $\mathbf{X}_{n1} \neq \boldsymbol{\theta}_0$ almost surely, which ensures that the $\mathbf{S}_{ni}$'s are well-defined with probability one. Finally, a sufficient (yet not necessary) condition for (iv) is that $\sqrt{n}\,\mathrm{E}[u_{n1}^2] \to \infty$ as $n \to \infty$. In other words, if (iv) does not hold, we must then have that, for some constant $C > 0$,

$$\mathrm{E}[(\mathbf{X}_{n1}'\boldsymbol{\theta}_0)^2] \geq 1 - \frac{C}{\sqrt{n}} \tag{3.3}$$

for infinitely many n. In the high-dimensional setup considered, (3.3) is extremely pathological, since it corresponds to the distribution of $\mathbf{X}_{n1}$ concentrating in *one* particular direction — namely, the direction $\boldsymbol{\theta}_0$ — in the expanding Euclidean space $\mathbb{R}^{p_n}$. Most importantly, it should

be noted that (ii) allows p_n to go to infinity in an arbitrary way with n, so that Theorem 3.1 provides a "(n,p)-universal" asymptotic distribution result for the modified Watson statistic.

The ratio decomposition of $\tilde{W}_n$ in (3.2) invites to base the proof of Theorem 3.1 on the Slutsky Lemma. The stochastic convergence of the denominator is taken care of in

Proposition 3.1. *Under the assumptions of Theorem 3.1,*

$$\frac{\frac{1}{n}\sum_{i=1}^n u_{ni}^2}{\mathrm{E}[u_{n1}^2]} \to 1$$

in quadratic mean as $n \to \infty$.

of Proposition 3.1. Since

$$\mathrm{E}\left[\left(\frac{\frac{1}{n}\sum_{i=1}^n u_{ni}^2}{\mathrm{E}[u_{n1}^2]} - 1\right)^2\right] = \frac{1}{(\mathrm{E}[u_{n1}^2])^2}\,\mathrm{E}\left[\left(\frac{1}{n}\sum_{i=1}^n u_{ni}^2 - \mathrm{E}[u_{n1}^2]\right)^2\right]$$
$$= \frac{1}{(\mathrm{E}[u_{n1}^2])^2}\,\mathrm{Var}\left[\frac{1}{n}\sum_{i=1}^n u_{ni}^2\right] = \frac{\mathrm{Var}[u_{n1}^2]}{n(\mathrm{E}[u_{n1}^2])^2} \le \frac{\mathrm{E}[u_{n1}^4]}{n(\mathrm{E}[u_{n1}^2])^2},$$

the result follows from Condition (iv) in Theorem 3.1. □

To establish Theorem 3.1, it is therefore sufficient to prove

Proposition 3.2. *Under the assumptions of Theorem 3.1,*

$$R_n := \frac{\sqrt{2(p_n-1)}}{n\mathrm{E}[u_{n1}^2]} \sum_{1\le i<j\le n} u_{ni}u_{nj}\mathbf{S}_{ni}'\mathbf{S}_{nj}$$

is asymptotically standard normal.

The proof of this proposition is much more delicate and will be based on the following martingale Central Limit Theorem; see Theorem 35.12 in [5].

Theorem 3.2. *Assume that, for each* n, $Z_{n1}, Z_{n2}, \ldots$ *is a martingale relative to the filtration* $\mathcal{F}_{n1}, \mathcal{F}_{n2}, \ldots$ *and define* $Y_{n\ell} = Z_{n\ell} - Z_{n,\ell-1}$. *Suppose that the* $Y_{n\ell}$*'s have finite second-order moments and let* $\sigma_{n\ell}^2 = \mathrm{E}[Y_{n\ell}^2 \,|\, \mathcal{F}_{n,\ell-1}]$ *(with* $\mathcal{F}_{n0} = \{\emptyset, \Omega\}$*). Assume that* $\sum_{\ell=1}^\infty Y_{n\ell}$ *and* $\sum_{\ell=1}^\infty \sigma_{n\ell}^2$ *converge with probability 1. Then, if, for* $n \to \infty$,

$$\sum_{\ell=1}^\infty \sigma_{n\ell}^2 = \sigma^2 + o_{\mathrm{P}}(1), \tag{3.4}$$

where σ *is a positive real number, and*

$$\sum_{\ell=1}^\infty \mathrm{E}\big[Y_{n\ell}^2\, \mathbb{I}[|Y_{n\ell}| \ge \varepsilon]\big] \to 0 \quad \forall \varepsilon > 0, \tag{3.5}$$

we have that $\sigma^{-1}\sum_{\ell=1}^{\infty} Y_{n\ell}$ *is asymptotically standard normal.*

In order to apply this result, we need to identify the distinct quantities in the present setting. Let $\mathcal{F}_{n\ell}$ be the σ-algebra generated by $\mathbf{X}_{n1},\ldots,\mathbf{X}_{n\ell}$ and denote by $\mathrm{E}_{n\ell}[.]$ the conditional expectation with respect to $\mathcal{F}_{n\ell}$. Then, letting

$$Y_{n\ell} := \mathrm{E}_{n\ell}[R_n] - \mathrm{E}_{n,\ell-1}[R_n] = \frac{\sqrt{2(p_n-1)}}{n\mathrm{E}[u_{n1}^2]} \sum_{i=1}^{\ell-1} u_{ni}u_{n\ell}\mathbf{S}_{ni}'\mathbf{S}_{n\ell}$$

for $\ell = 1,\ldots,n$ and (as in [5]) $Y_{n\ell} = 0$ for $\ell > n$, we clearly have that $R_n = \sum_{\ell=2}^{n} Y_{n\ell}$, where the $Y_{n\ell}$'s have finite second-order moments. Also, $\sum_{\ell=2}^{\infty} Y_{n\ell} = \sum_{\ell=2}^{n} Y_{n\ell}$ and $\sum_{\ell=2}^{\infty} \sigma_{n\ell}^2 = \sum_{\ell=2}^{n} \sigma_{n\ell}^2$, with $\sigma_{n\ell}^2 = \mathrm{E}_{n,\ell-1}[Y_{n\ell}^2]$ as in Theorem 3.2, and both converge with probability 1, as required. Now, the crucial conditions (3.4) and (3.5) are shown to hold in the subsequent lemmas (see the Appendix for the proofs).

Lemma 3.1. *Under the assumptions of Theorem 3.1,* $\sum_{\ell=2}^{n} \sigma_{n\ell}^2 \to 1$ *in quadratic mean as* $n \to \infty$.

Lemma 3.2. *Under the assumptions of Theorem 3.1,* $\sum_{\ell=2}^{n} \mathrm{E}[Y_{n\ell}^2\, \mathbb{I}[|Y_{n\ell}| > \varepsilon]] \to 0$ *as* $n \to \infty$ *for any* $\varepsilon > 0$.

These lemmas allow to use Theorem 3.2 to prove Proposition 3.2 which, jointly with Proposition 3.1, establishes Theorem 3.1. Clearly, the resulting high-dimensional Watson test, $\tilde{\phi}_n^W$, say, rejects the null hypothesis $\mathcal{H}_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$ in favor of $\mathcal{H}_1 : \boldsymbol{\theta} \neq \boldsymbol{\theta}_0$ at asymptotic level α whenever

$$\tilde{W}_n > \Phi^{-1}(1-\alpha),$$

where Φ denotes the cumulative distribution function of the standard normal distribution. As already pointed out when commenting the assumptions of Theorem 3.1, this test achieves asymptotic null size α irrespective of the way p_n goes to infinity with n.

For the problem considered above, [18] introduced the high-dimensional *sign* statistic

$$\tilde{S}_n := \frac{\sqrt{2(p_n-1)}}{n} \sum_{1\le i<j\le n} \mathbf{S}_{ni}'\mathbf{S}_{nj} \tag{3.6}$$

and showed that the (n,p)-universal asymptotic null distribution of $\tilde{S}_n$ is standard normal. In the next result, we identify assumptions on the sequence u_{n1} under which $\tilde{W}_n$ and $\tilde{S}_n$ are ((n,p)-universally) asymptotically equivalent in probability under the null.

Theorem 3.3. *Let the assumptions of Theorem 3.1 hold and further assume that (v)* $\mathrm{E}[u_{n1}^2]/(\mathrm{E}[u_{n1}])^2 \to 1$ *as* $n \to \infty$. *Then,* $\tilde{W}_n - \tilde{S}_n = o_{\mathrm{P}}(1)$ *as* $n \to \infty$.

of Theorem 3.3. Decompose $\tilde{W}_n - \tilde{S}_n$ into $A_n + B_n$, with

$$A_n = \left(\frac{\mathrm{E}[u_{n1}^2]}{\frac{1}{n}\sum_{i=1}^{n} u_{ni}^2} - 1 \right) \frac{\sqrt{2(p_n-1)}}{n\mathrm{E}[u_{n1}^2]} \sum_{1\le i<j\le n} u_{ni}u_{nj}\mathbf{S}_{ni}'\mathbf{S}_{nj}$$

and

$$B_n = \frac{\sqrt{2(p_n-1)}}{n} \sum_{1\le i<j\le n} \left(\frac{u_{ni}u_{nj}}{\mathrm{E}[u_{n1}^2]} - 1 \right) \mathbf{S}_{ni}'\mathbf{S}_{nj}.$$

Propositions 3.1 and 3.2 readily entail that $A_n = o_{\mathrm{P}}(1)$ as $n \to \infty$. As for B_n, we have (see the beginning of the Appendix for a recall on some results regarding expectations of the signs $\mathbf{S}_{ni}$)

$$\mathrm{E}[B_n^2] = \frac{2(p_n-1)}{n^2} \sum_{1\le i<j\le n} \mathrm{E}\left[\left(\frac{u_{ni}u_{nj}}{\mathrm{E}[u_{n1}^2]} - 1 \right)^2 (\mathbf{S}_{ni}'\mathbf{S}_{nj})^2 \right] = \frac{2}{n^2} \sum_{1\le i<j\le n} \mathrm{E}\left[\left(\frac{u_{ni}u_{nj}}{\mathrm{E}[u_{n1}^2]} - 1 \right)^2 \right]$$

$$= \frac{n-1}{n} \mathrm{E}\left[\left(\frac{u_{n1}u_{n2}}{\mathrm{E}[u_{n1}^2]} - 1 \right)^2 \right] = \frac{2(n-1)}{n} \mathrm{E}\left[1 - \frac{u_{n1}u_{n2}}{\mathrm{E}[u_{n1}^2]} \right] = \frac{2(n-1)}{n} \left(1 - \frac{(\mathrm{E}[u_{n1}])^2}{\mathrm{E}[u_{n1}^2]} \right),$$

which, in view of Condition (v), is $o(1)$ as $n \to \infty$. The result follows. □

This result shows that, quite intuitively, if u_{n1} becomes constant asymptotically (in the sense that $\mathrm{Var}[u_{n1}]/(\mathrm{E}[u_{n1}])^2 \to 0$), then the high-dimensional Watson test $\tilde{\phi}_n^W$ coincides with the sign test based on (3.6). This should be considered as the exception rather than the rule, though, since there is no particular reason why the distribution of $\mathbf{X}_{n1}$ should concentrate in (a possibly translated version of) the orthogonal complement of $\boldsymbol{\theta}_0$.

4 Spiked covariance matrices

Let $\mathbf{Y}_{n1}, \ldots, \mathbf{Y}_{nn}$ be a random sample from the p_n-dimensional multinormal distribution with mean zero and covariance matrix $\boldsymbol{\Sigma}$. For fixed $\boldsymbol{\theta}_0 \in \mathcal{S}^{p_n-1}$, we consider here the problem of testing the null hypothesis that $\boldsymbol{\Sigma}$ has a "$\boldsymbol{\theta}_0$-spiked" structure, that is, is of the form

$$\mathcal{H}_0^{\mathrm{spi}} : \boldsymbol{\Sigma} = \sigma^2(\mathbf{I}_{p_n} + \lambda\boldsymbol{\theta}_0\boldsymbol{\theta}_0'), \quad \text{for some } \sigma^2, \lambda > 0.$$

Consider the projections $\mathbf{X}_{ni} := \mathbf{Y}_{ni}/\|\mathbf{Y}_{ni}\|$, $i = 1, \ldots, n$, of the observations on the unit hypersphere, and let

$$\mathbf{S}_{ni} := \frac{\mathbf{X}_{ni} - (\mathbf{X}_{ni}'\boldsymbol{\theta}_0)\boldsymbol{\theta}_0}{\|\mathbf{X}_{ni} - (\mathbf{X}_{ni}'\boldsymbol{\theta}_0)\boldsymbol{\theta}_0\|}.$$

Under $\mathcal{H}_0^{\mathrm{spi}}$, (i) the $\mathbf{S}_{ni}$'s are mutually independent and are uniformly distributed over $\mathcal{S}^{p_n-1}(\boldsymbol{\theta}_0^{\perp}) := \{\mathbf{x} \in \mathcal{S}^{p_n-1} \,|\, \mathbf{x}'\boldsymbol{\theta}_0 = 0\}$; moreover, (ii) the $\mathbf{X}_{ni}'\boldsymbol{\theta}_0$'s are independent and identically distributed, and they are independent of the $\mathbf{S}_{ni}$'s. It is well-known that (i)-(ii) imply that the common distribution of the projected observations $\mathbf{X}_{ni}$ is rotationally symmetric about $\boldsymbol{\theta}_0$. Consequently, a high-dimensional test for $\boldsymbol{\theta}_0$-spikedness is the test, $\tilde{\phi}_n^{\mathrm{spi}}$ say, that rejects the null $\mathcal{H}_0^{\mathrm{spi}}$, at asymptotic level α, whenever

$$\tilde{W}_n^{\mathrm{spi}}(\mathbf{Y}_{n1}, \ldots, \mathbf{Y}_{nn}) := \tilde{W}_n(\mathbf{X}_{n1}, \ldots, \mathbf{X}_{nn}) > \Phi^{-1}(1-\alpha).$$

Theorem 3.1 ensures that $\tilde{\phi}_n^{\mathrm{spi}}$ has asymptotic null size α as soon as p_n goes to infinity with n (universal (n,p) asymptotics), which is illustrated in the simulations of the next section.

Typically, this test will show large powers against $\boldsymbol{\theta}$-spiked alternatives, with $\boldsymbol{\theta} \neq \boldsymbol{\theta}_0$.

5 Monte Carlo study

In this section, we conduct a Monte Carlo simulation study to check the validity of our universal asymptotic results related to both $\tilde{W}_n$ and $\tilde{W}_n^{\text{spi}}$. To do so, we generated, for every $(n,p) \in C \times C$, with $C = \{5, 30, 200, 1{,}000\}$, and for $\boldsymbol{\theta}_0$ the first vector of the canonical basis of $\mathbb{R}^p$, $M = 2{,}500$ independent random samples from each of the following p-dimensional distributions :

(i) the FvML distribution $\mathcal{R}(\boldsymbol{\theta}_0, F_{p,2})$ (see Section 2);

(ii) the Purkayastha distribution $\mathcal{R}(\boldsymbol{\theta}_0, G_{p,1})$, associated with $G_{p,\kappa}(t) = d_{p,\kappa} \int_{-1}^{t} (1 - s^2)^{(p-3)/2} \exp(-\kappa \arccos(s))\, ds$ $(t \in [-1,1])$, where $d_{p,\kappa}$ is a normalizing constant;

(iii) the multinormal distribution with mean zero and covariance matrix $\boldsymbol{\Sigma} = \mathbf{I}_p + (1/2)\boldsymbol{\theta}_0\boldsymbol{\theta}_0'$.

The modified Watson statistic $\tilde{W}_n$ was evaluated on the samples from (i)-(ii) (rotational symmetry about $\boldsymbol{\theta}_0$), while the statistic $\tilde{W}_n^{\text{spi}}$ was computed for each sample from (iii) ($\boldsymbol{\theta}_0$-spikedness). For each (n,p) regime considered, we report the corresponding histograms of $\tilde{W}_n$ and $\tilde{W}_n^{\text{spi}}$ in Figures 1-2 and in Figure 3, respectively (each histogram is based on $M = 2,500$ values of these statistics).

From Theorem 3.1 and the discussion in Section 4, histograms are expected to be approximately standard normal as soon as $\min(n,p)$ is large, in a universal way (that is, irrespective of the relative size of n and p). Inspection of the results shows that, for all three setups, the standard normal approximation is valid for moderate to large values of n and p, irrespective of the value of p/n, which confirms our universal asymptotic results. Note also that, for small p and moderate to large n (that is, $p = 5$ and $n \geq 30$), histograms are approximately (standardized) chi-square, which is consistent with classical fixed-p asymptotic results; see Section 2.

Acknowledgement

Christophe Ley thanks the Fonds National de la Recherche Scientifique, Communauté Française de Belgique, for support via a Mandat de Chargé de Recherche. Davy Paindaveine's research is supported by an A.R.C. contract from the Communauté Française de Belgique and by the IAP research network grant nr. P7/06 of the Belgian government (Belgian Science Policy).

Appendix: proofs of Lemmas 3.1 and 3.2

We recall that, under the assumptions of Theorem 3.1, the signs $\mathbf{S}_{ni}$ are uniformly distributed over $\mathcal{S}^{p_n-1}(\boldsymbol{\theta}_0^{\perp})$ (see Section 4) and that the u_{ni}'s are independent of the $\mathbf{S}_{ni}$'s, $i = 1, \ldots, n$. From Lemma A.1 in [18] it directly follows that, for fixed n, the quantities $\rho_{n,ij} := \mathbf{S}_{ni}'\mathbf{S}_{nj}$ are pairwise independent and satisfy $\mathrm{E}[\rho_{n,ij}] = 0$, $\mathrm{E}[\rho_{n,ij}^2] = 1/(p_n - 1)$, and $\mathrm{E}[\rho_{n,ij}^4] = 3/(p_n^2 - 1)$.

of Lemma 3.1. Rotational symmetry about $\boldsymbol{\theta}_0$ readily yields $\mathrm{E}[\mathbf{S}_{n\ell}\mathbf{S}_{n\ell}'] = \frac{1}{p_n-1}(\mathbf{I}_{p_n} - \boldsymbol{\theta}_0\boldsymbol{\theta}_0')$. The independence between the u_{ni}'s and $\mathbf{S}_{ni}$'s then provides

$$\sigma_{n\ell}^2 = \mathrm{E}_{n,\ell-1}[Y_{n\ell}^2] = \frac{2(p_n-1)}{n^2(\mathrm{E}[u_{n1}^2])^2}\sum_{i,j=1}^{\ell-1} u_{ni}u_{nj}\mathrm{E}[u_{n\ell}^2]\mathbf{S}_{ni}'\mathrm{E}[\mathbf{S}_{n\ell}\mathbf{S}_{n\ell}']\mathbf{S}_{nj} = \frac{2}{n^2\mathrm{E}[u_{n1}^2]}\sum_{i,j=1}^{\ell-1} u_{ni}u_{nj}\rho_{n,ij}.$$

Hence we obtain

$$\mathrm{E}\left[\sum_{\ell=2}^{n}\sigma_{n\ell}^2\right] = \frac{2}{n^2\mathrm{E}[u_{n1}^2]}\sum_{\ell=2}^{n}\sum_{i,j=1}^{\ell-1}\mathrm{E}[u_{ni}u_{nj}]\mathrm{E}[\rho_{n,ij}] = \frac{2}{n^2}\sum_{\ell=2}^{n}(\ell-1) = \frac{n-1}{n}. \tag{.7}$$

Moreover, the pairwise independence of the $\rho_{n,ij}$'s entails

$$\mathrm{Var}\left[\sum_{\ell=2}^{n}\sigma_{n\ell}^2\right] = \frac{4}{n^4(\mathrm{E}[u_{n1}^2])^2}\mathrm{Var}\left[\sum_{\ell=2}^{n}\sum_{i,j=1}^{\ell-1}u_{ni}u_{nj}\rho_{n,ij}\right] = \frac{4}{n^4(\mathrm{E}[u_{n1}^2])^2}\left\{T_1^{(n)} + 4\,T_2^{(n)}\right\},$$

with

$$T_1^{(n)} := \mathrm{Var}\left[\sum_{\ell=2}^{n}\sum_{i=1}^{\ell-1}u_{ni}^2\right] = \mathrm{Var}\left[\sum_{i=1}^{n-1}(n-i)u_{ni}^2\right] = \sum_{i=1}^{n-1}(n-i)^2\,\mathrm{Var}[u_{n1}^2] \le n^3\,\mathrm{Var}[u_{n1}^2]$$

and

$$\begin{aligned} T_2^{(n)} &:= \mathrm{Var}\left[\sum_{\ell=2}^{n}\sum_{1\le i<j\le \ell-1}u_{ni}u_{nj}\rho_{n,ij}\right] = \mathrm{Var}\left[\sum_{1\le i<j\le n-1}(n-j)u_{ni}u_{nj}\rho_{n,ij}\right] \\ &= \sum_{1\le i<j\le n-1}(n-j)^2\mathrm{Var}[u_{ni}u_{nj}\rho_{n,ij}] = \sum_{1\le i<j\le n-1}(n-j)^2\mathrm{E}[u_{ni}^2u_{nj}^2\rho_{n,ij}^2] \\ &= \frac{(\mathrm{E}[u_{n1}^2])^2}{p_n-1}\sum_{1\le i<j\le n-1}(n-j)^2 \le \frac{n^4(\mathrm{E}[u_{n1}^2])^2}{p_n-1}. \end{aligned}$$

Hence,

$$\mathrm{Var}\left[\sum_{\ell=2}^{n}\sigma_{n\ell}^2\right] \le \frac{4\mathrm{Var}[u_{n1}^2]}{n(\mathrm{E}[u_{n1}^2])^2} + \frac{16}{p_n-1} \le \frac{4\mathrm{E}[u_{n1}^4]}{n(\mathrm{E}[u_{n1}^2])^2} + \frac{16}{p_n-1} \to 0, \tag{.8}$$

in view of Conditions (ii) and (iv) from Theorem 3.1. Using (.7) and (.8) in

$$\mathrm{E}\left[\left(\sum_{\ell=2}^{n}\sigma_{n\ell}^2 - 1\right)^2\right] = \mathrm{Var}\left[\sum_{\ell=2}^{n}\sigma_{n\ell}^2\right] + \left(\mathrm{E}\left[\sum_{\ell=2}^{n}\sigma_{n\ell}^2 - 1\right]\right)^2$$

then establishes the result. □

of Lemma 3.2. Applying first the Cauchy-Schwarz inequality, then the Chebyshev inequality,

yields

$$\sum_{\ell=2}^{n} \mathrm{E}[Y_{n\ell}^2\, \mathbb{I}[|Y_{n\ell}| > \varepsilon]] \leq \sum_{\ell=2}^{n} \sqrt{\mathrm{E}[Y_{n\ell}^4]}\, \sqrt{\mathrm{P}[|Y_{n\ell}| > \varepsilon]} \leq \frac{1}{\varepsilon} \sum_{\ell=2}^{n} \sqrt{\mathrm{E}[Y_{n\ell}^4]}\, \sqrt{\mathrm{Var}[Y_{n\ell}]}.$$

Noting that $\mathrm{Var}[Y_{n\ell}] \leq \mathrm{E}[Y_{n\ell}^2] = 2(\ell-1)/n^2$, we obtain

$$\sum_{\ell=2}^{n} \mathrm{E}[Y_{n\ell}^2\, \mathbb{I}[|Y_{n\ell}| > \varepsilon]] \leq \frac{\sqrt{2}}{\varepsilon n} \sum_{\ell=2}^{n} \sqrt{\ell\, \mathrm{E}[Y_{n\ell}^4]}. \tag{.9}$$

Using the fact that $0 \leq u_{ni} \leq 1$ almost surely and the independence between the u_{ni}'s and the $\mathbf{S}_{ni}$'s, we get

$$\begin{aligned}
\mathrm{E}\left[\left(\sum_{i=1}^{\ell-1} u_{ni}u_{n\ell}\rho_{n,i\ell}\right)^4\right] &= \sum_{i,j,r,s=1}^{\ell-1} \mathrm{E}\big[u_{n\ell}^4 u_{ni}u_{nj}u_{nr}u_{ns}\rho_{n,i\ell}\rho_{n,j\ell}\rho_{n,r\ell}\rho_{n,s\ell}\big] \\
&= (\ell-1)(\mathrm{E}[u_{n1}^4])^2\mathrm{E}\big[\rho_{n,1\ell}^4\big] + 3(\ell-1)(\ell-2)\mathrm{E}[u_{n1}^4](\mathrm{E}[u_{n1}^2])^2\mathrm{E}\big[\rho_{n,1\ell}^2\rho_{n,2\ell}^2\big] \\
&= \frac{3(\ell-1)}{p_n^2-1}(\mathrm{E}[u_{n1}^4])^2 + \frac{3(\ell-1)(\ell-2)}{(p_n-1)^2}\mathrm{E}[u_{n1}^4](\mathrm{E}[u_{n1}^2])^2 \\
&\leq \frac{3}{(p_n-1)^2}\Big[\ell(\mathrm{E}[u_{n1}^4])^2 + \ell^2\mathrm{E}[u_{n1}^4](\mathrm{E}[u_{n1}^2])^2\Big]
\end{aligned}$$

which yields

$$\begin{aligned}
\mathrm{E}\big[Y_{n\ell}^4\big] &\leq \frac{4(p_n-1)^2}{n^4(\mathrm{E}[u_{n1}^2])^4} \times \frac{3}{(p_n-1)^2}\Big[\ell(\mathrm{E}[u_{n1}^4])^2 + \ell^2\mathrm{E}[u_{n1}^4](\mathrm{E}[u_{n1}^2])^2\Big] \\
&\leq \frac{12}{n^4}\left[\ell\,\frac{(\mathrm{E}[u_{n1}^4])^2}{(\mathrm{E}[u_{n1}^2])^4} + \ell^2\,\frac{\mathrm{E}[u_{n1}^4]}{(\mathrm{E}[u_{n1}^2])^2}\right].
\end{aligned}$$

Plugging into (.9), we conclude that

$$\begin{aligned}
\sum_{\ell=2}^{n} \mathrm{E}[Y_{n\ell}^2\, \mathbb{I}[|Y_{n\ell}| > \varepsilon]] &\leq \frac{\sqrt{24}}{\varepsilon n^3} \sum_{\ell=2}^{n} \sqrt{\ell^2\,\frac{(\mathrm{E}[u_{n1}^4])^2}{(\mathrm{E}[u_{n1}^2])^4} + \ell^3\,\frac{\mathrm{E}[u_{n1}^4]}{(\mathrm{E}[u_{n1}^2])^2}} \\
&\leq \frac{\sqrt{24}}{\varepsilon n^3} \sum_{\ell=2}^{n} \left(\ell\,\frac{\mathrm{E}[u_{n1}^4]}{(\mathrm{E}[u_{n1}^2])^2} + \ell^{3/2}\sqrt{\frac{\mathrm{E}[u_{n1}^4]}{(\mathrm{E}[u_{n1}^2])^2}}\right) \\
&\leq O(n^{-1})\,\frac{\mathrm{E}[u_{n1}^4]}{(\mathrm{E}[u_{n1}^2])^2} + O(n^{-1/2})\sqrt{\frac{\mathrm{E}[u_{n1}^4]}{(\mathrm{E}[u_{n1}^2])^2}},
\end{aligned}$$

which, in view of Condition (iv) from Theorem 3.1, is indeed $o(1)$. □

References

[1] Banerjee, A., Dhillon, I., Ghosh, J. & Sra, S. (2003). Generative model-based clustering of directional data. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, 19–28.

[2] Banerjee, A., Dhillon, I., Ghosh, J. & Sra, S. (2005). Clustering on the unit hypersphere using von mises-fisher distributions. *J. Mach. Learn. Res.* **6**, 1345–1382.

[3] Banerjee, A. & Ghosh, J. (2002). Frequency sensitive competitive learning for clustering on high-dimensional hyperspheres. In *Proceedings International Joint Conference on Neural Networks*, 1590–1595.

[4] Banerjee, A. & Ghosh, J. (2004). Frequency sensitive competitive learning for scalable balanced clustering on high-dimensional hyperspheres. *IEEE T. Neural Networ.* **15**, 702–719.

[5] Billingsley, P. (1995). *Probability and Measure*. New York, Chichester: Wiley, 3rd ed.

[6] Cai, T., Fan, J. & Jiang, T. (2013). Distributions of angles in random packing on spheres. *J. Mach. Learn. Res.* **14**, 1837–1864.

[7] Cai, T. & Jiang, T. (2012). Phase transition in limiting distributions of coherence of high-dimensional random matrices. *J. Multivariate Anal.* **107**, 24–39.

[8] Chen, S. & Qin, Y. (2010). A two-sample test for high-dimensional data with applications to gene-set testing. *Ann. Statist.* **38**, 808–835.

[9] Chen, S. X., Zhang, L.-X. & Zhong, P.-S. (2010). Tests for high-dimensional covariance matrices. *J. Amer. Statist. Assoc.* **105**, 810–819.

[10] Dryden, I. L. (2005). Statistical analysis on high-dimensional spheres and shape spaces. *Ann. Statist.* **33**, 1643–1665.

[11] Jiang, T. & Yang, F. (2013). Central limit theorems for classical likelihood ratio tests for high-dimensional normal distributions. *Ann. Statist.* **41**, 2029–2074.

[12] Johnstone, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *Ann. Statist.* **29**, 295–327.

[13] Ledoit, O. & Wolf, M. (2002). Some hypothesis tests for the covariance matrix when the dimension is large compared to the sample size. *Ann. Statist.* **30**, 1081–1102.

[14] Li, J. & Chen, S. X. (2012). Two sample tests for high-dimensional covariance matrices. *Ann. Statist.* **40**, 908–940.

[15] Mardia, K. V. & Jupp, P. E. (2000). *Directional Statistics*. John Wiley & Sons.

[16] Onatski, A., Moreira, M. & Hallin, M. (2013). Asymptotic power of sphericity tests for high-dimensional data. *Ann. Statist.* **41**, 1204–1231.

[17] Paindaveine, D. & Verdebout, T. (2013). Optimal rank-based tests for the location parameter of a rotationally symmetric distribution on the hypersphere. *ECARES Working Paper 2013-36* .

[18] Paindaveine, D. & Verdebout, T. (2013). Universal asymptotics for high-dimensional sign tests. *ECARES Working Paper 2013-40* .

[19] Saw, J. G. (1978). A family of distributions on the m-sphere and some hypothesis tests. *Biometrika* **65**, 69–73.

[20] Schott, J. (2007). Some high-dimensional tests for a one-way manova. *J. Multivariate Anal.* **98**, 1825–1839.

[21] Srivastava, M. S. & Fujikoshi, Y. (2006). Multivariate analysis of variance with fewer observations than the dimension. *J. Multivariate Anal.* **97**, 1927–1940.

[22] Srivastava, M. S., Katayama, S. & Kano, Y. (2013). A two sample test in high dimensional data. *J. Multivariate Anal.* **114**, 349–358.

[23] Srivastava, M. S. & Kubokawa, T. (2013). Tests for multivariate analysis of variance in high dimension under non-normality. *J. Multivariate Anal.* **115**, 204–216.

[24] Stam, A. J. (1982). Limit theorems for uniform distributions on spheres in high-dimensional euclidean spaces. *J. Appl. Probab.* **19**, 221–228.

[25] Watson, G. S. (1983). Limit theorems on high-dimensional spheres and stiefel manifolds. In *Studies in Econometrics, Time Series, and Multivariate Statistics*, S. Karlin, T. Amemiya & L. A. Goodman, eds. New York: Academic Press, 559–570.

[26] Watson, G. S. (1983). *Statistics on Spheres*. New York: Wiley.

[27] Watson, G. S. (1988). The langevin distribution on high dimensional spheres. *J. Appl. Statist.* **15**, 123–130.

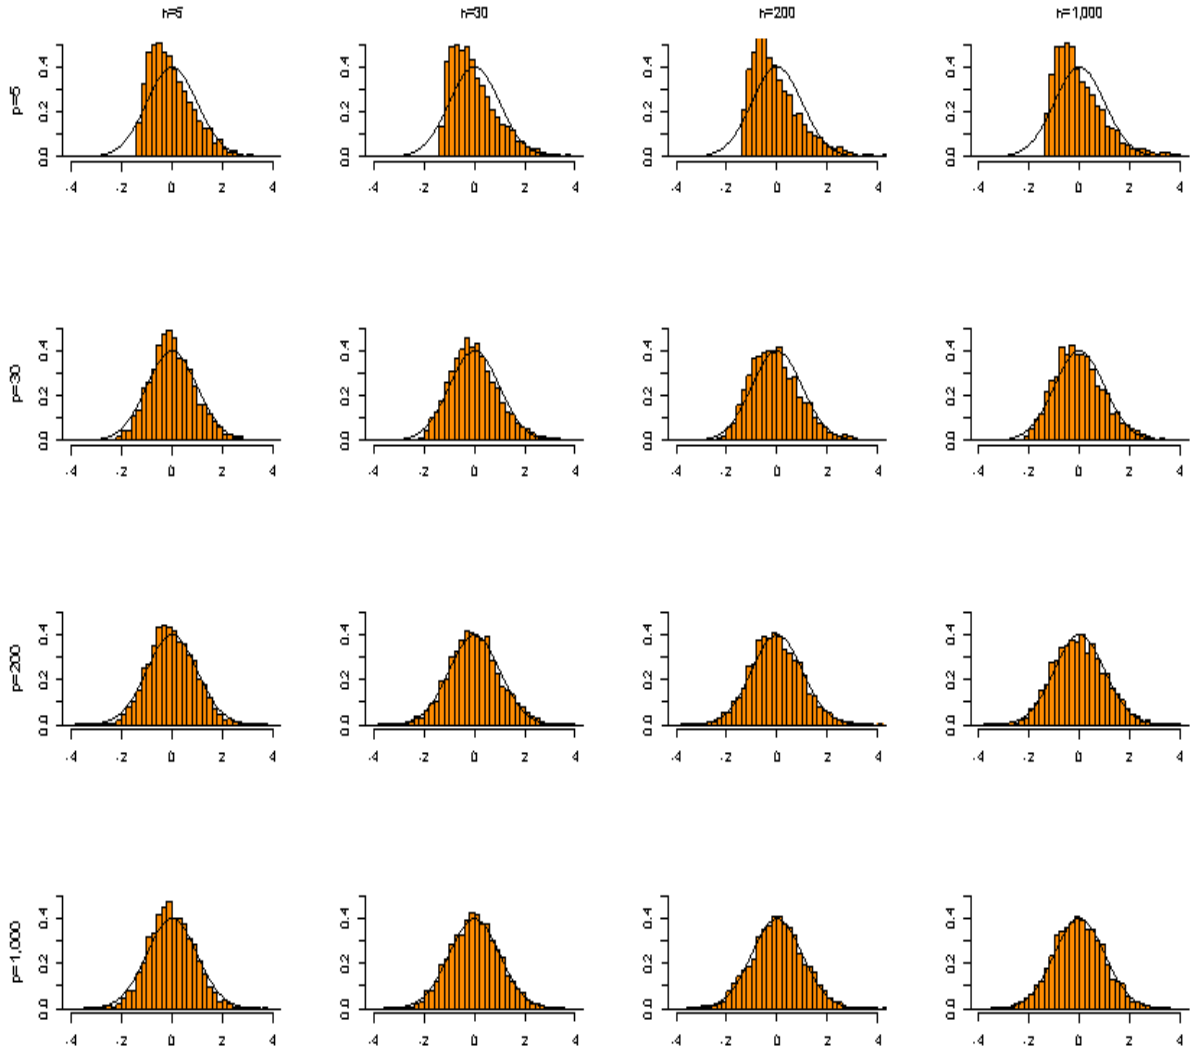

Figure 1: Histograms, for various values of n and p, of the modified Watson statistic $\tilde{W}_n$ evaluated on $M = 2,500$ random samples of size n from the p-dimensional FvML distribution with concentration $\kappa = 2$; see Section 5 for details.

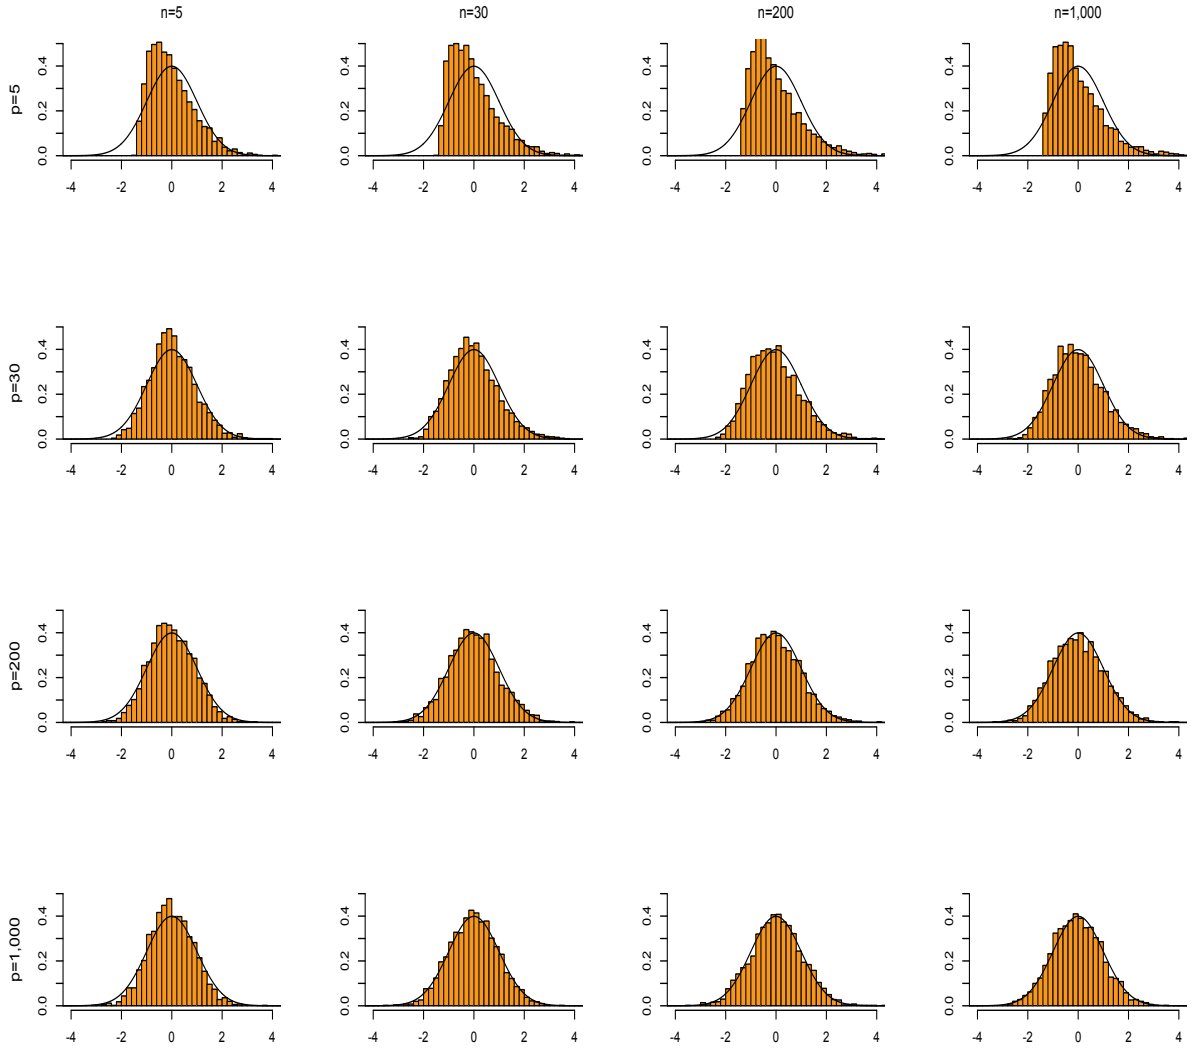

Figure 2: Histograms, for various values of n and p, of the modified Watson statistic $\tilde{W}_n$ evaluated on $M = 2,500$ random samples of size n from the p-dimensional Purkayastha distribution with concentration $\kappa = 1$; see Section 5 for details.

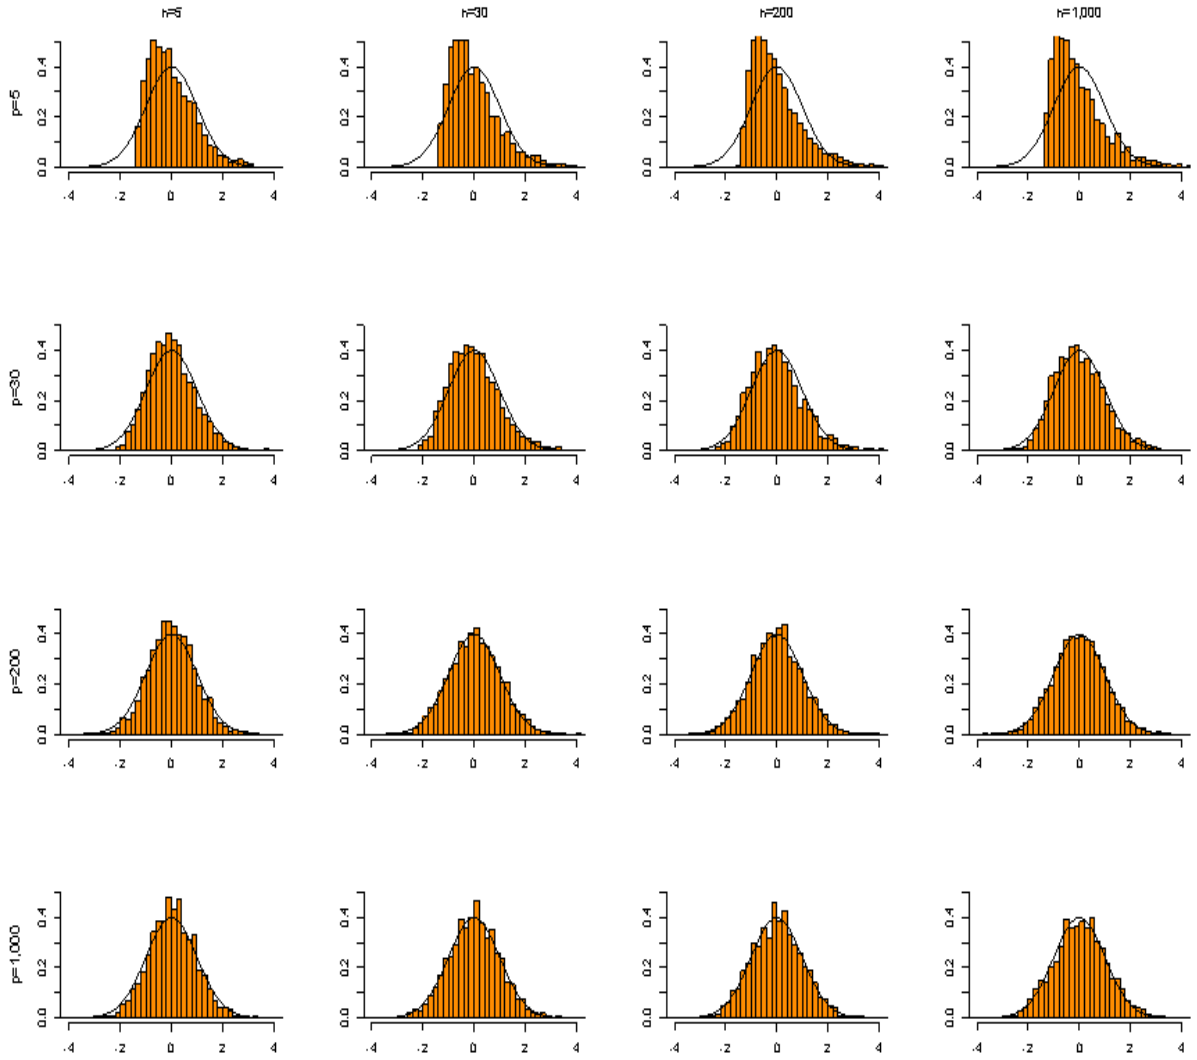

Figure 3: Histograms, for various values of n and p, of the test statistic $\tilde{W}_n^{\text{spi}}$ for $\boldsymbol{\theta}_0$-spikedness evaluated on $M = 2,500$ random samples of size n from the p-dimensional multinormal distribution with mean zero and covariance matrix $\boldsymbol{\Sigma} = \mathbf{I}_p + (1/2)\boldsymbol{\theta}_0\boldsymbol{\theta}_0'$; see Section 5 for details.